

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
APPLICATION FOR LETTERS PATENT

**Hierarchical, Probabilistic, Localized,
Semantic Image Classifier**

Inventor:
David Benedict Bradshaw

ATTORNEY's DOCKET NO. MS1-622US

1

TECHNICAL FIELD

2

3 This invention relates to semantic, content-based classification of localized
4 areas of an image as one of a number of multiple discriminating categories. More
5 particularly, this invention relates to one or more hierarchical, probabilistic
6 techniques for performing such semantic, content-based classification.

7

BACKGROUND

8

9 Digital images are increasingly more common as scanners and digital
10 cameras drop in price and increase in availability and function. As digital
11 photographers (amateurs and professionals alike) amass large collections of digital
12 photographs on their computers, the challenges involved with organizing,
13 querying, and accessing digital images grow.

14 Therefore, digital photographers need to utilize “image retrieval”
15 technology to accomplish their task. “Image retrieval” refers to a technology
16 focused on the organization of a library of digital images, the inquiry into such a
17 library, and the retrieval of selected images that meet the terms of such inquiry.

18

Content-Based Categorization Approaches

19

20 Images in a library may be organized and, thus, retrieved in an organized
21 fashion based upon their content. This may be accomplished manually or
22 automatically. An approach is manual if one or more persons largely perform the
23 analysis and categorization of images in the library and/or if one or more persons
24 largely browse the images to find those matching a query. An approach is

1 automated if these tasks are chiefly performed by a computer without significant
2 human interaction.

3 Content-based categorization and image retrieval approaches are beneficial
4 to both professionals and amateurs. Professional digital photographers (or any
5 other digital image acquirers) may have one or more extraordinarily large libraries
6 of images. They need to find the most cost-effective and efficient image retrieval
7 approach available. Similarly, amateur digital photographers wish to minimize
8 time spent organizing and retrieving images in their growing library of digital
9 images.

10
11 **Manual, Brute-Force, Content-Based Image Retrieval Approach**

12 A person may organize a small library of digital images by manually
13 analyzing, sorting, and categorizing each image into one or several labeled
14 computer file folders (i.e., directories). With this simple brute-force approach, a
15 person views and examines the content of each image. Based upon a manual
16 content-based analyses of an image, an analyzing person places the image into one
17 of several categories.

18 This technique has at least three major drawbacks:

- 19
- 20 • image classification is highly labor intensive;
 - 21 • image re-classification requires a manual re-examination of the
entire library; and
 - 22 • image querying is unavailable.

23 A re-classification requires a complete repetition of an already completely
24 manual and labor-intensive task. Furthermore, it is a time-consuming and a fully
25

1 manual task to locate a known image in such a library. A person must view and
2 examine every picture in the library or, at best, every picture in a given category.

3 To alleviate the problems of re-classification and image querying, some
4 conventional manual approaches associate relevant words with each image.
5 Therefore, the images may be classified and re-classified based upon their
6 associated “index words.” In addition, the images may be searched based upon
7 such words. However, querying and any subsequent re-classification are limited to
8 the set of words initially associated with images. Nevertheless, this approach still
9 requires a manual and time-consuming initial image analysis and word-
10 association.

11

12 Automated, Content-Based Image Retrieval Approach (CBIR)

13 Automated, Content-Based Image Retrieval (herein, simply referred to as
14 CBIR) has been actively researched in the last ten years because of the
15 increasingly large volumes of electronically stored information and the
16 corresponding requirement for high performance systems to access and manipulate
17 this information.

18 Initial conventional CBIR systems (classed herein as first-generation) give
19 a user (such as a professional or amateur user) direct access to the underlying low-
20 level feature space. Examples of such low-level feature spaces include those
21 corresponding to color and/or texture information. Examples of such systems
22 include the IBM Query By Image Content (QBIC) system (M. Flickner, et al.
Intelligent Multimedia Information Retrieval, p. 8-22. AAAI Press, 1997), the
23 Virage system (A. Gupta and R. Jain, *Communications of the ACM*, 40(5):71-79,
24 May 1997), and the VisualSEEk system (J. Smith and S. Chang. *Intelligent*
25

1 *Multimedia Information Retrieval*, p. 23-41. AAAI Press, 1997). Each of these
2 systems allows the user to specify a query in a number of ways based on the
3 features extracted by the system.

4 The user is also often allowed to specify how much weight to attach to each
5 of these features. However, it is now recognized in the field of image retrieval that
6 this explicit knowledge of the feature space does not help the user formulate a
7 query. Typically, the user simply cannot tell which low-level features, for a given
8 query, are more or less important.

9 There are a number of conventional CBIR systems (classified herein as
10 second-generation), which deliberately hide the low-level features from the user.
11 Instead of specifying texture and color combinations, the user supplies an example
12 image and asks for similar ones. This task is often called “Query By Example”
13 (QBE).

14 Although this alleviates the problem of knowing which low-level features
15 are important for a given query it immediately introduces another one. Namely, the
16 user must already have a good example of what the user wants prior to initiating
17 the query. Examples of such systems are the “texture of textures” system devised
18 by Debonet and Viola (J. S. DeBonet. Novel statistical multiresolution techniques
19 for image synthesis, discrimination, and recognition. Master’s thesis, M.I.T.
20 Learning and Vision Group, AI Lab, 1997) and the system developed by Rui et al.
21 at the University of Illinois.

22 **Image Feature Space**

24 Examples of an image feature space include a low-level image feature
25 space and a higher-level semantic image feature space. A point in an image

1 feature space may be referred to as a “feature”, a “feature vector”, or a “vector.”
2 Any image feature space of an image includes observable and measurable
3 characteristics of the content of such an image.

4 Low-level image feature spaces include basic observable and measurable
5 characteristics having little or no semantic value. Herein, a point within a low-
6 level image feature space may be referred to as a “low-level feature”, a “low-level
7 feature vector”, or a “low-level vector.”

8 Color and texture are examples of low-level features. Typically, the colors
9 and textures of an image’s content have little or no meaning themselves. Rather,
10 objects depicted in the image by a complex combination of colors and textures
11 have meaning. For example, a person is probably more interested in locating
12 images of their “home” than images of “tan-brown block with raised shadow
13 stripes at an angle of 20-55 degrees with a surrounding wash of green with spots
14 of brown.”

15 A higher-level semantic image feature space includes observable and
16 measurable characteristics with semantic value (i.e., meaning). Such higher-level
17 features may be constructed using supervised learning techniques (such as the
18 Kernel Fisher discriminant) in association with sets of extracted low-level
19 features. Herein, a point within a higher-level semantic image feature space may
20 also be referred to as a “higher-level feature”, a “higher-level feature vector”, a
21 “higher-level vector”, “semantic feature”, “semantic vector”, and the like.

22 Typically, conventional systems use low-level image feature space for
23 indexing.

1 **Semantic Content Extraction**

2 There are a number of publications that address the issue of extracting
3 semantic content from images. One of the first was Gorkani and Picard (Texture
4 orientation for sorting photos “at a glance”. In *Proc. of the IEEE Int. Conf. on*
5 *Pattern Recognition*, October 1994), who attempt to discriminate “city” from
6 “landscape” scenes using a texture orientation approach that is based on a
7 multi-scale steerable pyramid operating on 128×128 pixel blocks across an
8 image. Yiu (Image classifier using color cues and texture orientation. Master’s
9 thesis, Dept EECS, MIT, 1996) uses an identical texture extraction approach but
10 introduces color information to classify indoor and outdoor scenes. The procedure
11 used is based on a nearest neighbor approach combined with a support vector
12 machine classifier.

13 Gorkani and Picard use one block size (e.g., 128×128) and extract
14 orientations at multiple-frequency-scales to give them a discriminant. Their system
15 was developed for use on black & white photos; thus, it does not include color
16 information. Moreover, their system is non-probabilistic; rather, it uses a number
17 of hand-tuned threshold parameters to perform the classification.

18 Szummer and Picard (Indoor-outdoor image classifier. In *IEEE Int. Work.*
19 *on Content-based Access of Image and Vid. Databases*, January 1998) address the
20 same problem but combine a number of feature extraction techniques. An example
21 of such a combination of techniques includes a combination of Ohta color space
22 histograms (see Y. Ohta, T. Kanade, and T. Sakai. Color information for region
23 segmentation. *Computer Graphics and Image Processing*, 13:222-241, 1980) and
24 textural features based on a multi-resolution simultaneous autoregressive model
25

1 (see J. Mao and A. Jain. Texture classification and segmentation using
2 multiresolution simultaneous autoregressive models. *Pat. Rec.*, 25(2):173-188,
3 1992).

4 None of the conventional approaches described above are hierarchical.
5 Instead, they only analyze the image at one block size (the ‘multi-scale’
6 description indicates that the textures are extracted across a number of different
7 frequency scales). Moreover, none of these systems are probabilistic. To perform
8 the classification the first conventional system (Gorkani and Picard) uses a number
9 of hand-tuned threshold parameters, the second conventional system (Szummer
10 and Picard) uses a k-nearest neighbors approach, and the third conventional
11 system (Yiu) uses a support vector machine classifier.

12 In addition, Torralba and Oliva (A. B. Torralba and A. Oliva. Semantic
13 organization of scenes using discriminant structural templates. In *ICCV'99*, Corfu,
14 Greece, September 1999) address the issue of extracting semantic content from
15 images. They attempt to determine a set of real-valued ‘semantic axes’ in their
16 chosen feature space. While they assign real-values to each image in relation to
17 each semantic label rather than the more common binary classification approach,
18 they do not extend these real-values to a probabilistic representation. The three
19 axes chosen are: “Natural/Man-made”, “Open/Closed” (on the subset of natural
20 images) and “Expanded/Enclosed” (on the sub-set of man-made images). The
21 Torralba and Oliva system only outputs one real value per image which itself is
22 non-probabilistic.

1

Statistical Representation of Features

2 Within images, textures and other low-level features may be represented by
3 statistical properties. Such statistical properties are often used in the field of
4 object-recognition. The following publications provide good background in the
5 area of statistical representation of texture features:

- 6 • DeBonet and Viola (Multiresolution sampling procedure for analysis
7 and synthesis of texture images. In *SIGGRAPH'97*, pp. 361-368,
8 1997);
- 9 • Heeger and Bergen (Pyramid-based texture analysis/synthesis. In
10 *SIGGRAPH'95*, pp. 229-238, 1995);
- 11 • Zhu, Wu and Mumford (Filters, random fields and maximum
12 entropy. *Int. Journal of Computer Vision*, 27(2):1-20, March/April
13 1998);
- 14 • Wiskott et al. (Face recognition by elastic bunch graph matching. In
15 *IEEE Int. Conf. on Image Processing*, v. 1, 1997); and
- 16 • Rikert et al. (A cluster-based statistical model for object detection. In
17 *ICCV'99*, 1999).

18 The majority of the techniques covered in the above publications analyzes
19 and extract texture from an image over a number of scales at a given position in
20 the image.

1 Conventional Systems Using Semantic Representations

2 Presently, three conventional systems attempt to infer semantic
3 representations.

4 The first by Vailaya et al. categorizes ‘vacation images’ using a Bayesian
5 framework (A. Vailaya, M. Figueiredo, A. K. Jain, and HJ ZHang. Content-based
6 hierarchical classification of vacation images. In *IEEE Conf. on Multimedia
7 Computing and Systems*, v. 1, p. 518-523, 1999).

8 The second system by Fung and Loe is non-probabilistic, analyzing ‘natural
9 scenes’ (C. Y. Fung and K. F. Loe. A new approach for image classifier and
10 retrieval. In *22nd SIGIR Int. Conf. on Research and Development in Information
11 Retrieval*, p. 301-302, August 1999) and categorizing them into classes such as
12 “rural”, “forest”, “road” and “farm”.

13 The third system by Liu et al. classifies only medical images using a
14 probabilistic, feature selection method (Y. Liu, F. Dallaert and W.E. Rothfus.
15 Classification Driven Semantic Based Medical Image Indexing and Retrieval,
16 Technical Report: CMU-RI-TR-98-25, Carnegie Mellon University, 1998). Note
17 that the latter two are only applicable to restricted domains (the former being
18 ‘natural scenes’ and the latter being ‘medical images’).

19 The Vailaya et al. approach is a non-localized, non-hierarchical,
20 probabilistic system. There are at least three disadvantages of the Vailaya et al.
21 approach. Firstly, it requires thousands of images to train properly. This is because
22 their proposed system performs probability density estimation using a vector
23 quantization approach, which means that densities are estimated in the feature
24 vector space. If, for example, twelve dimensional feature vectors are used, Vailaya
25

et al. estimates the class conditional densities in that twelve dimensional space. Secondly, it only generates one value per image (thus, not localized results). Thirdly, the computational complexity of both training and testing is high. Vailaya et al. approach uses LUV and HSV color decompositions to extract color low-level features. It also uses multi-resolution, simultaneous auto-regressive approach to extract texture features.

The Fung and Loe approach is a non-probabilistic, non-localized, non-hierarchical system. A disadvantage of the Fung and Loe approach is that it obtains a single overall result for a given image by analyzing data extracted from 32x32 pixel sized patches. Another disadvantage of their approach is that they use a clustering, non-probabilistic approach to classify the patches that requires many thousands of images to train. A classification for the entire image is then determined using a singular value decomposition technique operating on a feature vector whose dimensions correspond to the results obtained from each of the patches. The Fung and Loe system uses an LUV color decomposition, extracting the mean and variance of these patches to use in the image feature space.

The Liu et al. method is a non-localized approach to classifying medical images of the brain. It is probabilistic, in nature, using a Bayesian approach to select which dimensions of the image feature space are best suited to a given semantic classification task. As mentioned earlier, the technique is domain specific in that it is only applicable in the medical imaging domain.

1 **SUMMARY**

2 Described herein is a technology for semantically classifying areas of an
3 image (and/or the images themselves) as one of a number of multiple
4 discriminating categories. More particularly, the technology employs one or more
5 hierarchical, probabilistic techniques for performing such classification. Such
6 technology is particularly useful in fields of image classification and image
7 retrieval.

8 The architecture of such technology employs multiple hierarchical layers.
9 The architecture is based on modeling class likelihoods at each of such layers
10 separately and then combining these to form an overall estimate of the posterior,
11 conditioned on the data. The task of combining the results across layers is made
12 more computationally efficient by assuming statistical independence between
13 layers.

14 This summary itself is not intended to limit the scope of this patent.
15 Moreover, the title of this patent is not intended to limit the scope of this patent.
16 For a better understanding of the present invention, please see the following
17 detailed description and appending claims, taken in conjunction with the
18 accompanying drawings. The scope of the present invention is pointed out in the
19 appending claims.

1 **BRIEF DESCRIPTION OF THE DRAWINGS**

2 The same numbers are used throughout the drawings to reference like
3 elements and features.

4 Fig. 1 is a two-part diagram illustrating the operation of the quad-tree and
5 the high-resolution approaches of an implementation of a hierarchical,
6 probabilistic, localized, semantic image classifier.

7 Fig. 2 shows two original images and the same images after
8 segmentationclassification performed by the quad-tree and the high-resolution
9 approaches of an implementation of a hierarchical, probabilistic, localized,
10 semantic image classifier.

11 Fig. 3 is a schematic block diagram showing an embodiment of a
12 hierarchical, probabilistic, localized, semantic image classifier.

13 Fig. 4 is a flow diagram showing a methodological implementation of a
14 hierarchical, probabilistic, localized, semantic image classifier.

15 Fig. 5 is an example of a computing operating environment capable of
16 implementing an implementation of a hierarchical, probabilistic, localized,
17 semantic image classifier.

18

19 **DETAILED DESCRIPTION**

20 The following description sets forth specific embodiments of the
21 hierarchical, probabilistic, localized, semantic image classifier that incorporate
22 elements recited in the appended claims. These embodiments are described with
23 specificity in order to meet statutory written description, enablement, and best-

1 mode requirements. However, the description itself is not intended to limit the
2 scope of this patent.

3 Described herein are one or more exemplary implementations of the
4 hierarchical, probabilistic, localized, semantic image classifier. The inventor
5 intends these exemplary implementations to be examples. The inventor does not
6 intend these exemplary implementations to limit the scope of the claimed present
7 invention. Rather, the inventor has contemplated that the claimed present invention
8 might also be embodied and implemented in other ways, in conjunction with other
9 present or future technologies.

10

11 **Introduction**

12 Described herein are one or more exemplary implementations of a
13 hierarchical, probabilistic, localized, semantic image classifier. One of such
14 exemplary implementations may be referred to as an exemplary “semantic
15 image classifier.” The exemplary semantic image classifier employs a probabilistic
16 approach to semantic labeling localized portions of an image.

17 With the semantic image classifier, the labels (i.e., classification or
18 categories) that are extracted from an image may be one of a number of multiple
19 discriminating categories. For example, areas of an image (or the images as a
20 whole) can be labeled as either “man-made” or “natural” objects. Alternative
21 categories may include “indoor” or “outdoor”; “sky” or “not sky”; “grass” or “not
22 grass”; and other such discriminating classifications. The classifications provided
23 herein are examples and not limitations on the classifications available.

24 The image classification is beneficial for indexing images based on their
25 semantic content. In other words, images are indexed based upon a higher-order

1 meaning of their content (e.g., “building”, “sky”, “indoor” or “outdoor”) rather
2 than a low-level observation of their content (e.g., “red”, “green”, “smooth
3 gradient”).

4 In the exemplary classifier described herein, a hierarchical framework is
5 utilized. Such a framework includes a hierarchy of levels (i.e., layers). At each of
6 the layers in a hierarchy, class likelihood probability densities are estimated. These
7 are then combined to form a posterior estimate of the probability of class
8 membership based on the evidence across layers.

9 The exemplary image classifier accurately generates localized semantic
10 labels (in a probabilistic setting). Specifically, the exemplary image classifier
11 addresses the problem of probabilistically labeling images, or parts of images, as
12 being made up of man-made (e.g. buildings, roads, cars) or natural (e.g. clouds,
13 forests, streams) objects. Furthermore, the exemplary classifier may address other
14 similar problems, such as the problem of determining whether an image was taken
15 inside or outside.

16 With the exemplary image classifier, semantics can be probabilistically
17 assigned to different areas of an image without having to resort to 3-D
18 reconstruction, point correspondences, or structure from motion techniques
19 normally associated with the conventional techniques, such as in the computer
20 vision community. Furthermore, the exemplary semantic image classifier does not
21 require a person to hand-segment hundreds of images on a pixel-by-pixel basis.

1 **Overview of the Exemplary Semantic Image Classifier**

2 Again, the exemplary semantic image classifier employs a multi-layer (i.e.,
3 hierarchy) architecture, which allows a localized approach to the semantic labeling
4 of images (or areas of an image).

5 To accomplish this labeling, the exemplary semantic image classifier
6 aggregates data over a number of different layers, each layer being a different
7 block size. The classifier combines the knowledge obtained from each of these
8 layers. This gives rise to a significantly more robust technique than a conventional
9 single layer approach.

10 With reference to Fig. 1, implementations of an exemplary image classifier
11 are described. Fig. 1 illustrates a quad-tree implementation (100) and a high-
12 resolution implementation (200). Fig. 1 includes a common image shown six times
13 in Fig. 1 as images 110, 120, 130, 210, 220, and 230. Each illustrated image in
14 Fig. 1 is identical to the others except for a block depicted thereon. Therefore,
15 references to the “common image” are to the same image shown six times (images
16 110, 120, 130, 210, 220, and 230) in Fig. 1. These images are from “Corel™
17 Gallery 1,000,000” collection.

18 The exemplary image classifier is called hierarchical because it divides and
19 analyses successive hierarchical layers (i.e., levels) of an image where each layer
20 consists of smaller blocks than the previous layer. For example, the exemplary
21 image classifier may use three layers.

22 At a top layer (L1 in Fig. 1), an image (such as image 110) is divided into
23 multiple large blocks (such as blocks 112). These blocks may be generically
24 called L1-blocks. Such blocks may be of a large size, such as 128 x 128 pixels.
25

1 Block 112a is shown where it appears on image 110 and in bold lines at 112a in
2 the hierarchical layer L1.

3 In the next layer (L2 in Fig. 1), one or more smaller blocks (such as block
4 122a of blocks 122) are coextensive with a portion of and associated with a large
5 block (such as block 112a of blocks 112) of the common image. These blocks may
6 be generically called L2-blocks. Such blocks are of a size smaller than L1-block.
7 For example, the L2 blocks may be 64 x 64 pixels while the L1-blocks were 128 x
8 128 pixels. Block 122a is shown where it appears on image 120 and in bold lines
9 at 122a in the hierarchical layer L2.

10 In the final layer (L3 in Fig. 1) of this example, one or more even smaller
11 blocks (such as block 132a of blocks 132) are coextensive with a portion of and
12 associated with a small block (such as block 122a of blocks 122) of the common
13 image. Therefore, these tiny blocks are coextensive with a portion of and
14 associated with a large block (such as block 112a of blocks 112) of the common
15 image. These tiny blocks may be generically called L3-blocks. Such blocks are of
16 a size smaller than L2-block. For example, the L3 blocks may be 32 x 32 pixels
17 while the L2-blocks were 64 x 64 pixels. Block 132a is shown where it appears
18 on image 130 and in bold lines at 132a in the hieratical layer L3.

19 Although not shown in Fig. 1, successively lower layers (L4, L5, etc.) may
20 continue forth with their accompanying smaller block sizes. The blocks of each
21 successively lower layers will be coextensive with a portion of and associated with
22 a block of the upper layers. After a hierarchy of layers is established, it is a
23 implementation choice regarding how many layers to employ.

24 The probabilistic analysis for each block is based upon extracting feature
25 vectors (such as color and texture) using known image processing (e.g., computer

1 vision, machine learning) techniques. The probabilistic determination for each
2 block is based upon a known classification technique called “Kernel Fisher’s
3 discriminant.”

4

5 Supervised Learning Techniques for Discriminating between Classes

6 Given a classification problem involving two or more classes (such as
7 Natural vs Man-made), Fisher’s linear discriminant aims to achieve an optimal
8 linear dimensionality reduction which then allows a discriminant to be determined.
9 The following reference describes Fisher’s linear discriminant: Bishop (Neural
10 Networks for Pattern Recognition, Oxford University Press, 1995). Although
11 Fisher’s linear discriminant works well for linearly separable problems it cannot
12 classify when the problem at hand is non-linearly separable.

13 The Kernel Fisher discriminant, although computationally more intensive,
14 can separate classes which are non-linearly separable and hence allows a much
15 wider range of problems to be solved. The following references describe Kernel
16 Fisher discriminant: Mika et al. (Fisher discriminant analysis with kernels, *Neural*
17 *Networks for Signal Processing IX*, pages 41-48, IEEE, 1999); and Roth et al.
18 (Nonlinear discriminant analysis using kernel functions, *Advances in Neural*
19 *Information Processing Systems 12* MIT Press, 2000).

20 Accordingly, the Kernel Fisher discriminant is utilized for implementations
21 of the exemplary image classifier described herein.

1 Assumption of Statistical Independence between Layers

2 It is assumed that the class likelihoods (i.e., probabilities) at each layer are
3 statistically independent of each other. Although there is an actual statistical
4 dependence because the blocks of each layer are contiguous and/or neighboring,
5 the exemplary image classifier assumes no dependence. This makes the
6 calculations easier and for a relatively small training data, produces better results.

7 One reason why this assumption is made in the exemplary image classifier
8 is that the data required to estimate the class likelihood need only be extracted
9 from homogeneous images. Herein, “homogeneous” refers to scenes that only
10 contain one class of image data (e.g., in the natural/man-made case, the image
11 consists of completely natural or completely man-made objects), whereas
12 “heterogeneous” refers to scenes containing both classes.

13 If statistical independence is not assumed, then the data used to estimate
14 class likelihood probability densities would have to be extracted from not only
15 homogeneous images but also heterogeneous images. The task of hand segmenting
16 hundreds of images on a pixel-by-pixel basis is extremely time-consuming
17 whereas selecting sets of images that completely consist of data from one class or
18 another is very simple. Heterogeneous examples are beneficial because otherwise
19 the extracted samples would only lie in those areas of the multidimensional feature
20 space corresponding to data that is homogeneous *across all layers*.

21 *The curse of dimensionality.* To generate the same density of examples in
22 the multidimensional feature space as that found in the one-dimensional case
23 requires an increase in the number of examples that is exponential with respect to
24 the increase in dimensionality. Given that 120 images are used in the exemplary

1 image classifier, the required increase in examples would be extremely large. In
2 practice, even a tenfold increase to 1200 images is extremely laborious to find and,
3 given the reason outlined above, even more difficult to hand-segment.

4 Those of ordinary skill in the art may have difficulty accepting this
5 assumption because it is an approximation. In reality, each layer is statistically
6 dependent upon the others because each layer includes some or all of the pixels
7 found in layers above and/or below in the hierarchy. Although this is an
8 approximation, it is a justified approximation.

9 Furthermore, it is assumed that a classified object is larger than the smallest
10 block size (or the lowest layer). Because of this implicit smoothness constraint,
11 the hierarchical approach of the exemplary semantic image classifier is unlikely to
12 correctly label particularly small objects of either class. A small object is one that
13 is smaller than the smallest block size (e.g., objects less than 16×16 pixels in
14 size).

15 Two Approaches: Quad-Tree and High-Resolution

16 Described herein are two approaches to image classification utilizing the
17 exemplary image classifier: “quad-tree” and “high-resolution” approaches.

18 Quad-Tree Approach

19 This approach is faster (than the high-resolution approach) because it
20 requires less computational processing because each pixel is analyzed only once
21 within each layer. However, the results are coarse and less detailed than the high-
22 resolution approach.

The quad-tree approach 100 is illustrated in Fig. 1. An image (such as the image shown at 110, 120, and 130) is divided into adjacent, tiled L1-blocks (such as blocks 112). Each L1-block (such as block 112a) is divided into adjacent, tiled L2-blocks (such as blocks 122). Each L2-block (such as block 122a) is further divided into adjacent, tiled L3-blocks (such as blocks 132). Although not shown in Fig. 1, each L3-block (such as 132a) is still further divided into adjacent, tiled blocks. And so forth.

In the quad-tree approach, each layer has a different block size. The block sizes decrease from L1 to L3. The blocks are arranged in a non-overlapping, quad-tree structure.

Table 1 defines four layers (L1-L4) and block sizes for an exemplary classifier using the quad-tree approach. For this example, all images used are either 256×384 pixels or 384×256 pixels in size, and have been extracted from the “CorelTM Gallery 1,000,000” collection.

Layer	Block Size	Blocks per image
L1	128×128	6
L2	64×64	24
L3	32×32	96
L4	16×16	384

Table 1: Definition of layers and block sizes.

This approach is called quad-tree because each subsequent layer of the hierarchy is equally divided into four (hence, “quad”) blocks. Thus, there are four branches from each layer. Of course, those of ordinary skill in the art understand

1 and appreciate that this is only one example of how to sub-divide a block at each
2 layer of a hierarchy. Those of ordinary skill in the art further understand and
3 appreciate that a block at a layer may be divided into any number of sub-blocks
4 (equally or unequally divided) for the next layer.

5

High-Resolution Approach

6

7 Unlike the quad-tree approach, there is one block per layer for each set of
8 associated blocks and each block is centered on a common pixel (i.e., point).
9 Therefore, each block at each layer are likely to overlap with one or more other
10 blocks at the same layer.

11 In the high-resolution approach, the exemplary image classifier analyzes
12 the image at a number of points, which are distributed across the image. Typically,
13 the points will be equally distributed across the image. These points may be called
14 “sampling points.”

15 The high-resolution approach 200 is illustrated in Fig. 1. A series of
16 overlapping L1-blocks (an example of one L1-block is 212) is formed over an
17 image (such as the image shown at 210, 220, and 230). Each L1-block is centered
18 on its own sampling point (i.e., pixel). Image 210 shows the location and relative
19 size of L1-block 212. By way of example, an L1-block may be 128 x 128 pixels.

20 Each L1-block overlaps its neighboring L1-blocks by a given number of
21 pixels. That number is an implementation choice to balance redundant processing
22 with smoother results. The degree of overlap is directly proportional to the
23 increase in computation because of reprocessing common pixels in multiple L1-
24 blocks.

1 Each L1-block (such as block 212) has one L2-block (such as block 222)
2 associated therewith. Each L2-block is centered on the same sampling point as its
3 associated L1-block (shown in dashed lines as block 212a). Image 220 shows the
4 location and relative size of L2-block 222. By way of example, an L2-block may
5 be 64 x 64 pixels.

6 Similarly, Each L2-block (such as block 222) has one L3-block (such as
7 block 232) associated therewith. Each L3-block is centered on the same sampling
8 point as its associated L2-block, which in turn is centered on the same sampling
9 point as its associated L1-block (shown in dashed lines as block 212a). Image 230
10 shows the location and relative size of L3-block 232. By way of example, an L3-
11 block may be 32 x 32 pixels.

12 This pattern continues for blocks at subsequently lower layers until the
13 lowest layer and smallest block is reach. The choice of the lowest layer and
14 smallest block is a design choice for each implementation. Nevertheless, all of the
15 layered blocks in a set of hierarchical associated blocks are centered on a common
16 pixel, this being one of the sampling points mentioned previously. Moreover, for
17 each sampling point, there is a set of hierarchical associated blocks.

18
19 Training

20 For both approaches, in order to obtain probabilistic labels, probability
21 densities representing the class likelihoods are estimated at each layer. In other
22 words, a set of training images is used to calculate class likelihoods.

23 To achieve this, a set of feature vectors is obtained for each class (for
24 example, the natural/man-made categorization may be considered) at each layer
25 (i.e., each block size) for images in the training set. This is done by selecting a set

1 of natural and a set of man-made homogeneous images from which feature vectors
2 can be extracted. A different and appropriate training set is used for different
3 discriminating categories (such as inside/outside). Having obtained a set of
4 example feature vectors for each class at each block size the kernel Fisher's
5 discriminant method is used which then allows class conditional densities to be
6 estimated. This is explained later, in the 'Probabilistic Labeling Technique'
7 section.

8

9 Both Approaches

10 With both approaches, all associated blocks are analyzed in the same
11 manner. More specifically, each block is analyzed to probabilistically determine
12 whether the image data in a block contains one of two discriminating categories of
13 objects (e.g., either man-made object or natural object). The probabilities obtained
14 from the associated blocks are then combined to form an estimate of the posterior
15 probability of class membership. This posterior estimate is compared to a given
16 threshold to determine which of the two discriminating categories the image data
17 corresponds to. This is explained later, in the 'Probabilistic Labeling Technique'
18 section. In the case of the quad-tree approach, this result is used to label the
19 smallest block (e.g., the block at L3). In the case of the high-resolution approach,
20 it is used to label the area surrounding the common pixel (i.e., sample point) on
21 which the associated blocks are centered.

22 Fig. 2 illustrates the results of each approach. Fig. 2 shows two different
23 original images A (300) and B (400). This shows the segmentation of the images
24 into either the man-made or natural class.

1 310 and 410 illustrate the segmentation results of the exemplary image
2 classifier using the high-resolution approach for images A and B, respectively. In
3 310 and 410, the darkened areas are classified as natural and the remaining areas
4 are classified as man-made.

5 320 and 420 illustrate the segmentation results of the exemplary image
6 classifier using the quad-tree approach for images A and B, respectively. In 320
7 and 420, the darkened areas are classified as natural and the remaining areas are
8 classified as man-made.

9 Fig. 2 illustrates that the high-resolution approach generally provides
10 smoother and more accurate segmentations than the quad-tree approach. Although
11 it cannot be shown in the results, the quad-tree approach is computationally more
12 efficient than the high-resolution approach.

13 **Extracting Feature Vectors**

15 The feature vectors have twenty-six dimensions. To make the task of
16 estimating probability densities easier, dimensionality reduction is performed and
17 then class likelihood densities estimated. Once class probability distributions have
18 been determined, labeling of an image can be performed, first by extracting the
19 image feature vectors from all blocks at all layers. Class likelihoods are then
20 estimated allowing posterior probabilities, conditioned on the data across layers, to
21 be determined by using Bayes rule.

22 There are two primary feature vectors: color and texture. These low-level
23 features are extracted from images. These features are utilized by the exemplary
24 image classifier to produce semantic level features.

1 The following discussion focuses on extracting feature vectors in the
2 context of the quad-tree approach. However, those of ordinary skill in the art
3 understand that the high-resolution approach uses the same or similar feature
4 extraction techniques.

5 The domain of the image is denoted as D and the area over which a block
6 extends is denoted as B where $B \subset D$. The quad-tree parent of this block is now
7 defined as $Q(B)$. The quad-tree decomposition is recursive in nature such that at a
8 given layer l (the letter L) in the decomposition:

$$9 \quad Q_l(B) = Q(Q_{l+1}(B)) \quad (1)$$

10 To aid clarity, in the following text a block at layer l is denoted as B_l where:

$$12 \quad B_l = Q_l(B) \quad (2)$$

14 This decomposition approach can be seen in the quad-tree approach 100 of
15 Fig. 1. The index l in the quad-tree definition corresponds to the layer (L1, L2,
16 L3, etc.) in the hierarchy as defined in Table 1. At each layer, feature vectors from
17 blocks B_1, B_2, B_3, \dots are extracted. After dimensionality reduction, the likelihood of
18 class membership can be estimated.

19 Color Extraction

21 There are many different models available for low-level color extraction.
22 The primary aim when extracting color is to obtain a set of values, which are as
23 de-correlated from each other as possible. Two common models used to achieve
24 this are the Ohta color model (see Ohta referenced in the background section) and
25 the HSV color model (Hue, Saturation and Value). On comparing these methods, it

is found that both the Ohta and HSV components had very similar correlation properties with both sets of components being significantly more de-correlated than those of the RGB (Red, Green, Blue) model did. Although they both give very similar results, the Ohta decomposition (i.e., transformation) is more easily computed than the HSV transformation, and so the former is used with the exemplary semantic image classifier. However, those skilled in the art understand that the latter or any other such color model may be employed in an implementation of an semantic image classifier.

Subsequently, the three components of the Ohta color transformation corresponding to image I are denoted as I_{o1} , I_{o2} , and I_{o3} . The first of these corresponds to the luminance information; the latter two correspond to the chrominance information.

The chrominance information corresponding to block B_1 is determined as follows:

$$C_1(B_1) = \int_{B1} I_{o2}(r) dr \quad (3)$$

$$C_2(B_1) = \int_{B1} I_{o3}(r) dr \quad (4)$$

where r denotes a particular pixel position in the image.

Texture Extraction

The low-level texture extraction approach is based on the complex wavelet transform (CWT). The CWT developed by Kingsbury (The dual-tree complex wavelet transform: A new efficient tool for image restoration and enhancement. In *EUSIPCO'98*, v. 1, pp. 319-322. EURASIP, 1998) is an efficient way of implementing a set of critically sampled Gabor-like wavelets. Gabor

wavelets/filters have been used by those investigating both semantic content and classification problems. They have also been shown to be similar in nature to the function of simple cells in the primary visual cortex of primates. With the exemplary image classifier, the CWT is used rather than Gabor wavelets because of the significantly reduced computational load.

Gabor-like impulse responses provide directional selectivity with six orientations at each scale. Although the filters may be approximately orthogonal, they do not generally provide statistically independent outputs.

To succinctly represent the proposed procedure, the following notation is introduced:

- the wavelet function at scale s , and orientation θ is denoted as ϕ_s^θ ;
- the orientation can take one of six values $\theta = \{15^\circ, 45^\circ, -75^\circ, -45^\circ, -15^\circ\}$;
- the θ in the following text refers to an index into the vector (e.g., $\theta \in I : \{1\dots6\}$); and
- the “ $*$ ” symbol denotes the convolution operator.

The response across the luminance image I_{o1} , extracted using the Ohta transformation, to each of the wavelet functions is determined as follows:

$$I_s^\theta = I_{o1} * \phi_s^\theta \quad (5)$$

The $o1$ is dropped to aid clarity. The energy response to the wavelet function at scale s , and orientation θ is defined for a block B_1 as:

$$(6)$$

$$T_s^\theta(B_1) = \int_{B1} (I_s^\theta(r))^2 dr$$

For the exemplary image classifier described herein, wavelet functions corresponding to scales 1-4 are used with six orientations at each scales giving rise to twenty-four texture based features per feature vector.

Thus, the feature vectors of the exemplary image classifier have twenty-six dimensions, twenty-four texture features and two color features. Using the terms defined above, the feature vector at a particular block B and layer l is found by concatenating the texture based features with the color based features as follows:

$$\mathbf{X}_l^B = [T_1^1(B_1), T_1^2(B_1), \dots, T_4^6(B_1), C_1(B_1), C_2(B_1)] \quad (7)$$

Probabilistic Labeling Technique

With this exemplary image classifier, there are two classes, which are denoted as C_k where $k \in 1,2$. For example, in the natural/manmade example, C_1 corresponds to natural and C_2 corresponds to man-made. Probabilities are denoted as $P(\bullet)$ and probability densities as $p(\bullet)$. To determine posterior estimates of class membership Bayes rule is used. This requires prior probabilities of class membership (denoted as $P(C_k)$), a way of estimating the likelihood of class membership for a given block B at a given layer l (denoted as $P(X_l^B | C_k)$) and a way of combining these likelihood estimates from a number of layers. Estimating the priors, $P(C_k)$, presents no problem. If there is no knowledge about the images to be analyzed, they can be set at 0.5.

1 Estimating the Class Likelihoods

2 To determine an estimate of the likelihood of class membership, estimation
3 of class conditional densities for each layer l must be undertaken these densities
4 being denoted as $p(X_l^B | C_k)$. To achieve this density estimation, dimensionality
5 reduction is undertaken through the use of the Kernel Fisher's discriminant
6 technique.

7 This kernel method initially transforms the feature vectors from layer l into
8 a high dimensional space via a non-linear mapping. It is this non-linear mapping
9 that makes this invention suitable for classifying non-linearly separable data (e.g.
10 sky vs. not sky or face vs. not face). In this invention a Gaussian Radial Basis
11 Function kernel is used to achieve the non-linear mapping, of course, those of
12 ordinary skill in the art understand and appreciate that this is only one choice of a
13 number of possible non-linear mappings.

14 Once in this high dimensional space a vector is determined in a similar
15 manner to that used in the Fisher's discriminant method (hence the name); this
16 vector is denoted subsequently as p_l , this discriminative approach gives the "best"
17 vector p_l onto which to project feature vectors in the sense that the projection
18 maximizes inter-class separation whilst minimizing intra-class distance. This
19 method is based on a number of assumptions regarding the distribution of the
20 feature vectors of each class that those of ordinary skill in the art understand and
21 appreciate.

22 Dimensionality reduction is performed by projecting the feature vector
23 extracted from block B at layer l (denoted as X_l^B) onto p_l , thus resulting in a scalar
24 value (i.e. the data is one dimensional). This value is subsequently denoted as X_l^B .

1 This makes probability density estimation very easy but throws away information
2 that could otherwise be possibly used to discriminate between two classes.

3 Having determined sets of projected values from the feature vectors from
4 each class, probability densities for the class likelihoods can be estimated. The
5 likelihoods are modeled using normal distributions because when the
6 dimensionality reduction step is undertaken it approximates to summing a set of
7 independent, random variables and thus the central limit theorem can be invoked.
8 The mean and variance of the one-dimensional normal distributions are then found
9 using the maximum likelihood approach.

10

11 Combining Class Likelihoods

12 As discussed earlier, it is assumed that the likelihoods at each layer are
13 statistically independent of each other when estimating the probability of class
14 membership given data at a number of layers.

15 In the following discussion, the block index B is dropped to aid clarity.
16 Given that data has been extracted from a number of block sizes, 1 to L , at a given
17 sampling point the posterior conditioned on this data is:

$$18 \quad P(C_k | X_1, \dots, X_L) = \frac{P(X_1, \dots, X_L | C_k)}{P(X_1, \dots, X_L)} P(C_k) \quad (8)$$

$$19 \quad \approx \frac{\prod_{l=1}^L P(X_l | C_k)}{P(X_1, \dots, X_L)} P(C_k) \quad (9)$$

20
21 The denominator can be evaluated using the chain rule as follows:
22

$$23 \quad P(X_1, \dots, X_L) = P(X_L | X_{L-1}, \dots, X_1) \dots P(X_2 | X_1) P(X_1) \quad (10)$$

1
2 Each factor in this expansion of the denominator can be evaluated by
3 marginalizing over the class variable C_k :

4
5 (11)
6
$$P(X_l | X_{l-1}, \dots, X_1) = \sum_{k=1,2} P(X_l | C_k) P(C_k | X_{l-1}, \dots, X_1)$$

7 Evaluating Equation 10 in a top down, recursive manner: At the top layer
8 (e.g., L1), the class likelihood estimates can be evaluated and the posterior
9 conditioned on the data X_1 determined. These results can then be used at layer two
10 (e.g., L2), combined with the class likelihood estimates at this layer to evaluate the
11 relevant factor in the expansion of the denominator as described in Equation 10.
12 Now, the posterior conditioned on the data $\{X_1, X_2\}$ can be determined. This
13 process is repeated until the lowest layer L (e.g., L4) is reached.

14 A quad-tree decomposition implies that at each level in the hierarchy, the
15 number of samples across an image increases by a factor of four. Estimation of
16 class likelihoods occurring between sampling points is achieved by using bilinear
17 interpolation based on the nearest four sample points available.

18 Note that in the case considered in these example problems are 2-class
19 problems. This results in the following equation:

20
21 (13)
22
$$P(C_1 | X_1, \dots, X_L) = 1 - P(C_2 | X_1, \dots, X_L)$$

23 To determine whether the sample under consideration belongs to class 1 or
24 class 2 equation 9 is evaluated, if the result is > 0.5 the sample is assigned as
25

1 belonging to class 1 otherwise it is assigned as belonging to class 2 (as illustrated
2 by equation 13).

3 Although equation 13 is directed to a 2-class problem, multi-class problems
4 may be addressed using the approach described herein.

5

More on the High Resolution Approach

6

7 The high-resolution approach uses a different sampling strategy than the
8 quad-tree approach. Specifically, in the high-resolution approach, the blocks
9 contributing to a sampling point are all centered at the same pixel position. In
10 using this approach, samples can be extracted at any position across the image. An
11 exemplary “grid” may be set to have 4x4 spacing. Unlike the quad-tree approach,
12 blocks in the same layer from adjacent samples will overlap.

13 Although the high-resolution approach is computationally less efficient
14 than the quad-tree approach, it generally gives superior segmentations to the
15 quad-tree approach.

16

Exemplary Semantic Image Classification System

17

18 Fig. 3 shows a semantic image classification system 500, which is an
19 embodiment of the exemplary semantic image classifier. This embodiment may be
20 implemented in software, hardware, or a combination thereof.

21 This embodiment includes an image retriever 510 for retrieving images
22 from an image library, such as library 552, or some other image source. Once a
23 subject image is chosen, a hierarchy definer 520 selects sampling points across the
24 image. Typically, such sampling points are equally distributed across the image.
25 The definer 520 defines L1-blocks centered on each point. For each of such L1-

blocks, it also defines associated hierarchical layers with one or more blocks per layer. The block(s) at each successively lower layer is smaller than and at least partially coextensive with one or more blocks in the layer immediately above.

For the quad-tree approach, each layer (except at L1) is composed of four adjacent, tiled blocks. In effect, the tiled blocks are formed by quartering the block of the layer immediately above. See the quad-tree approach (100) illustrated in Fig. 1.

For the high-resolution approach, each layer is centered on a common sampling point and is composed of only one block. Each block (except at L1) is smaller than the block of the layer immediately above. See the high-resolution approach (200) illustrated in Fig. 1.

After the blocks and the layered hierarchy sets have been defined, a block analyzer 530 analyzes the block(s) of each layer of each set to probabilistically determine whether the image data in a block contains one of two discriminating categories of objects. Such analysis includes the following (which are described in detail above): extracting feature vectors and estimating class likelihoods.

After such analysis, a probability combiner 540 combines the class likelihoods (probabilities) of hierarchical associated blocks. This is done for each set of associated blocks.

After that, an image segment classifier 550 compares the combined probability using equation 9 to determine which of the two discriminating categories the image data of a block corresponds to. The classifier 550 does this for each hierarchical set. The classifier 550 classifies blocks based upon the results of such a comparison. For example, a 16 x 16 block at layers L4 may be labeled

1 “man-made” or “natural” because the combined probability was above (or below)
2 a given threshold.

3 Furthermore, the classifier 550 may display the results on display 554
4 and/or save its results to the image library 552.

5 For the quad-tree approach, the estimate of class likelihoods occurring
6 between sampling points is achieved by using bilinear interpolation based on the
7 nearest group of sample points available.

8 When this system processes an entire image, segments (i.e., portions, areas,
9 etc.) of the image will be labeled one of multiple discriminating categories. Fig. 2
10 shows examples of such segmentation.

11 **Methodological Implementation of the Exemplary Semantic Image Classifier**

12 Fig. 4 shows a methodological implementation of the exemplary semantic
13 image classifier performed by the semantic image classification system 500. This
14 methodological implementation may be performed in software, hardware, or a
15 combination thereof.

16 At 610 of Fig. 4, the exemplary semantic image classifier retrieves an
17 image from an image library (or from any other source). At 612, it selects
18 sampling points across the image. Typically, such sampling points are equally
19 distributed across the image.

20 At 614, the exemplary image classifier forms L1-blocks centered on each
21 point. For each of such L1-blocks, it also defines associated hierarchical layers
22 with one or more blocks per layer. The block(s) at each successively lower layer
23 is smaller than and at least partially coextensive with one or more blocks in the
24 layer immediately above.

1 For the quad-tree approach, each layer (except at L1) is composed of four
2 adjacent, tiled blocks. In effect, the tiled blocks are formed by quartering the
3 block of the layer immediately above. See the quad-tree approach (100) illustrated
4 in Fig. 1.

5 For the high-resolution approach, each layer is centered on a common
6 sampling point and is composed of only one block. Each block (except at L1) is
7 smaller than the block of the layer immediately above. See the high-resolution
8 approach (200) illustrated in Fig. 1.

9 At 616 of Fig. 4, the exemplary image classifier analyzes the block(s) of
10 each layer of each set to probabilistically determine whether the image data in a
11 block contains one of two discriminating categories of objects. Such analysis
12 includes the following (which are described in detail above): extracting feature
13 vectors (such as color and texture) and estimating class likelihoods.

14 At 618, the class likelihoods (probabilities) of hierarchical associated
15 blocks are combined. This is done for each set of associated blocks. At 620, for
16 each hierarchical set, the exemplary image classifier compares the combined
17 probability to a given threshold to determine to which of the two discriminating
18 categories that the image data of a block corresponds. In the quad-tree case the
19 estimate of class likelihoods occurring between sampling points is achieved by
20 using bilinear interpolation based on the nearest group of sample points available.

21 At 622 of Fig. 4, the exemplary image classifier classifies blocks based
22 upon the results of such comparison. For example, a 16 x 16 block at layers L4
23 may be labeled “man-made” or “natural” because the combined probability was
24 above (or below) a given threshold.

When this methodological implementation is applied to an entire image, segments (i.e., portions, areas, etc.) of the image will be labeled one of multiple discriminating categories. Fig. 2 shows examples of such segmentation.

Exemplary Computing Environment

Fig. 5 illustrates an example of a suitable computing environment 920 on which an exemplary image classifier may be implemented.

Exemplary computing environment 920 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of an exemplary image classifier. Neither should the computing environment 920 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary computing environment 920.

The exemplary image classifier is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for use with an exemplary image classifier include, but are not limited to, personal computers, server computers, thin clients, thick clients, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

An exemplary image classifier may be described in the general context of computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects,

1 components, data structures, etc. that perform particular tasks or implement
2 particular abstract data types. An exemplary image classifier may also be
3 practiced in distributed computing environments where tasks are performed by
4 remote processing devices that are linked through a communications network. In
5 a distributed computing environment, program modules may be located in both
6 local and remote computer storage media including memory storage devices.

7 As shown in Fig. 5, the computing environment 920 includes a general-
8 purpose computing device in the form of a computer 930. The components of
9 computer 920 may include, by are not limited to, one or more processors or
10 processing units 932, a system memory 934, and a bus 936 that couples various
11 system components including the system memory 934 to the processor 932.

12 Bus 936 represents one or more of any of several types of bus structures,
13 including a memory bus or memory controller, a peripheral bus, an accelerated
14 graphics port, and a processor or local bus using any of a variety of bus
15 architectures. By way of example, and not limitation, such architectures include
16 Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA)
17 bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA)
18 local bus, and Peripheral Component Interconnects (PCI) bus also known as
19 Mezzanine bus.

20 Computer 930 typically includes a variety of computer readable media.
21 Such media may be any available media that is accessible by computer 930, and it
22 includes both volatile and non-volatile media, removable and non-removable
23 media.

24 In Fig. 5, the system memory includes computer readable media in the form
25 of volatile memory, such as random access memory (RAM) 940, and/or non-

1 volatile memory, such as read only memory (ROM) 938. A basic input/output
2 system (BIOS) 942, containing the basic routines that help to transfer information
3 between elements within computer 930, such as during start-up, is stored in ROM
4 938. RAM 940 typically contains data and/or program modules that are
5 immediately accessible to and/or presently be operated on by processor 932.

6 Computer 930 may further include other removable/non-removable,
7 volatile/non-volatile computer storage media. By way of example only, Fig. 5
8 illustrates a hard disk drive 944 for reading from and writing to a non-removable,
9 non-volatile magnetic media (not shown and typically called a “hard drive”), a
10 magnetic disk drive 946 for reading from and writing to a removable, non-volatile
11 magnetic disk 948 (e.g., a “floppy disk”), and an optical disk drive 950 for reading
12 from or writing to a removable, non-volatile optical disk 952 such as a CD-ROM,
13 DVD-ROM or other optical media. The hard disk drive 944, magnetic disk drive
14 946, and optical disk drive 950 are each connected to bus 936 by one or more
15 interfaces 954.

16 The drives and their associated computer-readable media provide
17 nonvolatile storage of computer readable instructions, data structures, program
18 modules, and other data for computer 930. Although the exemplary environment
19 described herein employs a hard disk, a removable magnetic disk 948 and a
20 removable optical disk 952, it should be appreciated by those skilled in the art that
21 other types of computer readable media which can store data that is accessible by a
22 computer, such as magnetic cassettes, flash memory cards, digital video disks,
23 random access memories (RAMs), read only memories (ROM), and the like, may
24 also be used in the exemplary operating environment.
25

1 A number of program modules may be stored on the hard disk, magnetic
2 disk 948, optical disk 952, ROM 938, or RAM 940, including, by way of example,
3 and not limitation, an operating system 958, one or more application programs
4 960, other program modules 962, and program data 964.

5 Each of such operating system 958, one or more application programs 960,
6 other program modules 962, and program data 964 (or some combination thereof)
7 may include an embodiment of an exemplary image classifier. More specifically,
8 each may include an embodiment of image retriever, hierarchy definer, block
9 analyzer, probabilities determiner, probability combiner, image segment classifier.

10 A user may enter commands and information into computer 930 through
11 input devices such as keyboard 966 and pointing device 968 (such as a “mouse”).
12 Other input devices (not shown) may include a microphone, joystick, game pad,
13 satellite dish, serial port, scanner, or the like. These and other input devices are
14 connected to the processing unit 932 through a user input interface 970 that is
15 coupled to bus 936, but may be connected by other interface and bus structures,
16 such as a parallel port, game port, or a universal serial bus (USB).

17 A monitor 972 or other type of display device is also connected to bus 936
18 via an interface, such as a video adapter 974. In addition to the monitor, personal
19 computers typically include other peripheral output devices (not shown), such as
20 speakers and printers, which may be connected through output peripheral interface
21 975.

22 Computer 930 may operate in a networked environment using logical
23 connections to one or more remote computers, such as a remote computer 982.
24 Remote computer 982 may include many or all of the elements and features
25 described herein relative to computer 930.

1 Logical connections shown in Fig. 5 are a local area network (LAN) 977
2 and a general wide area network (WAN) 979. Such networking environments are
3 commonplace in offices, enterprise-wide computer networks, intranets, and the
4 Internet.

5 When used in a LAN networking environment, the computer 930 is
6 connected to LAN 977 via network interface or adapter 986. When used in a
7 WAN networking environment, the computer typically includes a modem 978 or
8 other means for establishing communications over the WAN 979. The modem
9 978, which may be internal or external, may be connected to the system bus 936
10 via the user input interface 970 or other appropriate mechanism.

11 Depicted in Fig. 5, is a specific implementation of a WAN via the Internet.
12 Computer 930 typically includes a modem 978 or other means for establishing
13 communications over the Internet 980. Modem 978, which may be internal or
14 external, is connected to bus 936 via interface 970.

15 In a networked environment, program modules depicted relative to the
16 personal computer 930, or portions thereof, may be stored in a remote memory
17 storage device. By way of example, and not limitation, Fig. 5 illustrates remote
18 application programs 989 as residing on a memory device of remote computer
19 982. It will be appreciated that the network connections shown and described are
20 exemplary and other means of establishing a communications link between the
21 computers may be used.

22 **Exemplary Operating Environment**

23 Fig. 5 illustrates an example of a suitable operating environment 920 in
24 which an exemplary image classifier may be implemented. Specifically, the
25

1 exemplary image classifier(s) described herein is implemented (wholly or in part)
2 by any program module 960-962 and/or operating system 958 in Fig. 5 or a
3 portion thereof.

4 The operating environment is only an example of a suitable operating
5 environment and is not intended to suggest any limitation as to the scope or use of
6 functionality of the exemplary image classifier(s) described herein. Other well
7 known computing systems, environments, and/or configurations that may be
8 suitable for use with an exemplary image classifier include, but are not limited to,
9 personal computers (PCs), server computers, hand-held or laptop devices,
10 multiprocessor systems, microprocessor-based systems, programmable consumer
11 electronics, wireless phones and equipments, general- and special-purpose
12 appliances, application-specific integrated circuits (ASICs), network PCs,
13 minicomputers, mainframe computers, distributed computing environments that
14 include any of the above systems or devices, and the like.

15

Computer-Executable Instructions

16

17 An implementation of an exemplary image classifier may be described in
18 the general context of computer-executable instructions, such as program modules,
19 executed by one or more computers or other devices. Generally, program modules
20 include routines, programs, objects, components, data structures, etc. that perform
21 particular tasks or implement particular abstract data types. Typically, the
22 functionality of the program modules may be combined or distributed as desired in
23 various embodiments.

1 **Computer Readable Media**

2 An implementation of an exemplary image classifier may be stored on or
3 transmitted across some form of computer readable media. Computer readable
4 media can be any available media that can be accessed by a computer. By way of
5 example, and not limitation, computer readable media may comprise “computer
6 storage media” and “communications media.”

7 “Computer storage media” include volatile and non-volatile, removable and
8 non-removable media implemented in any method or technology for storage of
9 information such as computer readable instructions, data structures, program
10 modules, or other data. Computer storage media includes, but is not limited to,
11 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
12 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic
13 tape, magnetic disk storage or other magnetic storage devices, or any other
14 medium which can be used to store the desired information and which can be
15 accessed by a computer.

16 “Communication media” typically embodies computer readable
17 instructions, data structures, program modules, or other data in a modulated data
18 signal, such as carrier wave or other transport mechanism. Communication media
19 also includes any information delivery media.

20 The term “modulated data signal” means a signal that has one or more of its
21 characteristics set or changed in such a manner as to encode information in the
22 signal. By way of example, and not limitation, communication media includes
23 wired media such as a wired network or direct-wired connection, and wireless

1 media such as acoustic, RF, infrared, and other wireless media. Combinations of
2 any of the above are also included within the scope of computer readable media.

3

4 **Conclusion**

5 Although the hierarchical, probabilistic, localized, semantic image classifier
6 has been described in language specific to structural features and/or
7 methodological steps, it is to be understood that the hierarchical, probabilistic,
8 localized, semantic image classifier defined in the appended claims is not
9 necessarily limited to the specific features or steps described. Rather, the specific
10 features and steps are disclosed as preferred forms of implementing the claimed
11 present invention.

12
13
14
15
16
17
18
19
20
21
22
23
24
25